Reduced classification error probability with ternary correlation filters

John D. Downie

NASA Ames Research Center
Moffett Field, CA. 94035

ABSTRACT

Much of the filter design work that has been performed to date for filter SLMs with both constrained and unconstrained modulation characteristics has been concerned with optimizing the design for certain performance criteria associated only with the correlation function of the target image. However, in most likely application scenarios there will be multiple objects that may populate the field of view, and the most important correlation performance criterion is ultimately the probability of correct classification of a given object as either belonging to the in-class set or the out-of-class set. In this work, we study the problem of designing ternary phase and amplitude filters (TPAFs) that reduce the probability of image misclassification. We use the Fisher ratio as a measure of the correct classification rate, and we attempt to maximize this quantity in our filter designs. Given the non-analytical nature of the design problem, we employ a simulated annealing optimization technique. We present computer simulation results for several cases including single in-class and out-of-class image sets and multiple image sets corresponding to the design of synthetic discriminant function filters. We find significant reductions in expected rates of classification error in comparison to BPOFs and other TPAF designs.

1. INTRODUCTION

Binary phase-only filters (BPOFs) and ternary phase and amplitude filters (TPAFs) have merited study for implementation in optical correlators because they can be easily encoded on current spatial light modulator (SLM) devices. High-speed programmable SLMs offer the promise of real-time pattern recognition and classification for applications such as target tracking and autonomous robotic vision systems. Much of the filter design work that has been performed to date for filter SLMs with both constrained and unconstrained modulation characteristics has been concerned with optimizing the design for certain performance criteria associated only with the correlation function of the target image. Examples of such performance measures include the signal-to-noise ratio[1] (SNR), the peak-to-correlation energy[2,3] (PCE), and the Horner efficiency[4]. However, in most likely application scenarios there will be multiple objects that may populate the field of view, and the most important and basic correlation performance criterion will be the probability of correctly classifying a given object as either belonging to the in-class set or the out-of-class set. That is, a successful filter design must be able to clearly discriminate between the desired in-class image set and an out-of-class set, even in the presence of image noise.

In this work, we study the problem of designing ternary phase and amplitude filters that reduce the probability of image misclassification. Recent work has addressed the issue of optimally disciminating between a target image and white background noise[5], but here we are concerned with the situation of two or more distinct image classes, and the need to distinguish between them with a given filter. This situation has been addressed recently in the context of fully complex filters[6,7], but our study is aimed at the specific type of ternary encoded filter that is implementable on the magneto-optic spatial light modulator. The Fisher ratio (FR) is a quantitative measure that is directly related to the correct classification rate, and we attempt to maximize this quantity to obtain better filter designs in

0-8194-1275-9/93/$6.00

comparison to the conventional BPOF and TPAF designs that maximize another criterion such as SNR. The specific design algorithm followed here was to first design a BPOF to match the image or images defined as the in-class set, and then optimize the non-zero region of support (ROS) of the filter to improve the probability of correct classification between the in-class set, and another pre-defined out-of class set. White Gaussian image noise was assumed to be corrupting the images. We used a simulated annealing optimization technique to find the filter region of support, which in principle allows one to find the globally optimal solution while avoiding sub-optimal local solutions. We present computer simulation results for several cases including single in-class and out-of-class image sets and multiple image sets corresponding to the design of synthetic discriminant function filters. We find significant reductions in expected rates of classification error for ternary filters designed to maximize FR with a simulated annealing algorithm in comparison to BPOFs and other TPAF designs.

2. FILTER DESIGN ALGORITHM USING SIMULATED ANNEALING

To define the problem, consider Figure 1 which shows the probability density functions of both the in-class and out-of-class image correlation peaks for a hypothetical pair of image sets. They are illustrated as Gaussian distributions centered around their respective mean values. Even though we will assume that the image noise is white noise with each pixel having a Gaussian distribution with a mean of zero, this is not necessary for the peak distributions to be Gaussian. The central limit theorem predicts this result since the central correlation peak error is the sum of a large number of independent random variables (as many as there are image pixels). For zero-mean white noise, the means of the two distributions are equal to the peak values in the absence of noise, while the two standard deviations are equal and are determined by the standard deviation of the image noise distribution.

Referring again to Figure 1, it is easy to show that in order to equalize the error probabilites of the two variables, the correlation threshold t is located at the mean of the expected values of the in-class and out-of-class peaks,

$$t = \frac{1}{2}\left(E[c_{ic}] + E[c_{oc}]\right) \tag{1}$$

where c_{ic} is the in-class peak value and c_{oc} is the the out-of-class peak value. Correlation peaks measured higher than t are classified as in-class images while those lower than t are classified as out-of-class. If the variances of the two variables were unequal, the optimal threshold t would shift either to the right or left. Given the Gaussian nature of the two distributions and our assumption that their standard deviations are equal, we can write the probability of correct identification as the probability that the in-class peak will be greater than t, or

$$P_{ID} = \int_{t}^{\infty} \frac{1}{\sqrt{2\pi}\,\sigma_c} \exp \frac{-\left(x - E[c_{ic}]\right)^2}{2\,\sigma_c^{\,2}}\, dx \tag{2}$$

where σ_c is the standard deviation of the in-class (and out-of-class) correlation peaks. We can further simplify the expression for P_{ID} by letting z represent a normalized Gaussian random variable with mean 0 and variance 1. Then it is easy to show that

$$P_{ID} = \int_{-\infty}^{\left(E[c_{ic}] - t\right)/\sigma_c} \frac{1}{\sqrt{2\pi}} \exp \frac{-z^2}{2}\, dz \tag{3}$$

which can be evaluated easily by numerical integration or reference to printed tables of this function. Looking at the upper limit of the integration, we can write

$$a \equiv \frac{E[c_{ic}] - t}{\sigma_c}$$

$$= \frac{1}{2} \frac{E[c_{ic}] - E[c_{oc}]}{\sigma_c} \quad (4)$$

and note that P_{ID} is a monotonically increasing function of a. This analysis of the probability of correct identification is the basis of the Fisher ratio[8] performance measure which is defined as

$$FR = \frac{| E[c_{ic}] - E[c_{oc}] |^2}{\frac{1}{2}[\sigma_{ic}^2 + \sigma_{oc}^2]} \quad (5)$$

and it is thus clear that for $\sigma_{ic} = \sigma_{oc} = \sigma_c$,

$$FR = 4\, a^2 \quad (6)$$

as they are defined here. Since P_{ID} is an increasing function of a, it is also an increasing function of FR. We may thus use either a or FR as the basis of our filter design and attempt to maximize either measure to obtain the optimal filter.

In the work presented here, we used the measure FR as defined in Eq. 5 to assess the performance of a given filter. As implementable on the magneto-optic device[9], a ternary filter H(u,v) has possible pixel encoding values of (-1, 0, +1). Rather than attempt to find the most general ternary filter design that maximizes FR, we limit our algorithm to finding the best region-of-support (ROS) function for a binary phase-only filter. The ROS is defined as those filter pixels having a non-zero value. The BPOF function for a given image s(x,y) is

$$B(u,v) = \begin{cases} +1 & \text{real}(S(u,v)\, e^{-j\beta}) > 0 \\ -1 & \text{real}(S(u,v)\, e^{-j\beta}) < 0 \end{cases} \quad (7)$$

where S(u,v) is the Fourier transform of s(x,y) and in theory one must search through β-space from 0 to π in order to maximize the correlation peak. In practice, the optimal value of β is usually 0, and thus the BPOF B(u,v) can be expressed as the sign of the real part of the image Fourier transform. The filter H(u,v) is written as

$$H(u,v) = M(u,v) \cdot B(u,v) \quad (8)$$

where M(u,v) is a binary mask function having pixel values of 0 and 1 that manifests the ROS. It is this function M(u,v) that we seek to optimize in our design in order to maximize the FR measure. In the most general optimization we would allow each pixel to assume any of the 3 state values in the design process. However, this would significantly lengthen the design time in comparison to our approach, so we have effectively traded some optimality in performance for a design that can be obtained in a reasonable time. Also, an advantage of optimizing the ROS only is that the phase of each pixel within the ROS remains matched to the input image through the expression in Eq. 7.

Given the ternary encoding of the desired filter, and an assumption of zero-mean white image noise, it is easy to show that the expression for FR can be re-written as

$$FR = \frac{\left[\,|c_{ic}| - |c_{oc}|\,\right]^2}{\sigma_n^2 \sum_{u,v} M(u,v)} \quad (9)$$

where the denominator is simply the variance of the image noise multiplied by the number of pixels in the filter ROS, and c_{ic} and c_{oc} represent the correlation peak values in the absence of image noise. The form of Eq. 9, because of the fact that the out-of-class training image correlation peaks may be anywhere in the correlation plane, is such that precludes an analytic design of H(u,v). Thus we turn to an iterative optimization technique such as simulated annealing. Simulated annealing (SA) is an algorithm that is often used to solve complex optimization problems and in theory is able to find the globally optimal solution. SA has been previously applied to the design of BPOF correlation filters in order to increase the discrimination between two similar images with success[10]. Our effort here extends that work in the sense that we are applying the technique to ternary encoded filters (TPAFs) but we are optimizing the Fisher ratio as a direct measure of the probability of correct identification rather than the ratio of correlation peak strengths, which does not account for performance in the presence of image noise. We also apply the technique to designing synthetic discriminant filters (SDFs) which allow a degree of image distortion over a specified range.

The basis of the SA algorithm used is as follows: a perturbation is made to the system in the form of a change in state of the mask function M(u,v) at a given pixel. The new value of FR is computed after the perturbation, and the change in FR at the i^{th} iteration is expressed as

$$\Delta FR = FR^{i} - FR^{i-1} . \tag{10}$$

If ΔFR is positive, then the perturbation is always accepted. If the change in Fisher ratio is negative, corresponding to a decrease in the filter performance, then the perturbation is accepted with some probability. The probability is expressed as

$$P[\Delta FR] = \exp\left(\frac{\Delta FR}{kT}\right) \tag{11}$$

where k is a constant and T is the so-called temperature parameter. Since ΔFR as used in Eq. 10 is negative, the probability is a number between 0 and 1. In order to model the physical process of annealing, the temperature T is initially high so that there is a high probability of acceptance even when ΔFR is negative with a large magnitude. This helps to prevent the optimization from getting stuck in local minima. Gradually the temperature is cooled, so that eventually T is very low, and there is a very low probability of acceptance of changes that produce a negative ΔFR. Thus the algorithm should gradually "settle" into the global optimum. The cooling schedule of T is set by the algorithm designer. In this study, we used the schedule

$$T(j) = \alpha^{r} \cdot T(0) \tag{12}$$

where j is the iteration number, T(0) is the intial temperature, and r is given by

$$r = \text{floor}\left(\frac{j}{C}\right). \tag{13}$$

The variables T(0), C, and α are set at the beginning of the optimization. In Eq. 13, C is a number that specifies the number of iterations spent at the current temperature before T is decreased by the multiplying factor α.

3. SIMULATIONS AND RESULTS

3.1 Single in-class and out-of-class images

In our initial experiment with the use of the SA technique to maximize the Fisher ratio performance metric, we address the simplest case of a single in-class image and a single out-of-class image. We wish to maximize the probability of correct identification of a given input image that is a

noisy version of either one or the other, based on correlation with the filter under design. The images used for this simulation are shown in Figure 2. The BPOF function B(u,v) was first calculated for the in-class image as given in Eq. 7 with $\beta = 0$. The SA algorithm employed was designed to scan through the rows of the filter mask function M(u,v) and at each row, a column was picked at random. The value of the mask M(u,v) at the pixel defined by this row and column was then switched to its opposite state (0 or 1), thus creating a perturbation to the state of the filter. The effect of the perturbation was analyzed by calculating the new value of FR and comparing it to the previous value of FR. If the change was positive, the perturbation was automatically accepted. If it was negative, a probability was calculated according to Eq. 11, and a random number generated between 0 and 1 was compared to this probability to decide whether or not the change was accepted. This process continued with the temperature T changing according to Eq. 12 until the algorithm converged to a solution. The initial value of T was set such that even the largest negative values of ΔFR were given a high probability of being accepted in the beginning of the process. By the end, T was low enough that in effect only positive changes to FR were accepted.

Because the calculation of FR involves the determination of 2 correlation peak values, the SA design process could be quite slow if we computed by brute force the correlation function of both the in-class and out-of-class images after each filter perturbation. Thus we significantly speeded up the process by considering only the changes to the correlation functions effected by the filter perturbation. For the in-class image, we assumed that the correlation peak was at the center of the field (for a centered image), and thus we were concerned with only a single pixel. The central peak is expressed as

$$c_{ic} = \sum_{u,v} S_{ic}(u,v) \cdot H(u,v) = \sum_{u,v} S_{ic}(u,v) \cdot M(u,v)\, B(u,v) \tag{14}$$

and thus we simply add or subtract the quantity $S_{ic}(u,v)B(u,v)$ when the mask function at (u,v) is changed to either 1 or 0, respectively. The out-of-class correlation is not quite as easy, however, because in this case the correlation peak value can occur anywhere in the correlation field. As a compromise between speed and accuracy, we therefore assumed that the out-of-class peak would be found in the central 20x20 pixel region (in the 64x64 correlation plane) and updated only the values in that region in a manner similar to the in-class correlation central pixel, taking into account the appropriate complex values of the Fourier kernel at each given correlation plane position. We then took the maximum value in that 20x20 pixel region as the current value of c_{oc}. We found this technique to greatly speed up the design process with little limitation to the generality of the solution.

The results of the design of an FR-maximized TPAF for the single in-class and out-of-class images in Figure 2 are shown in Figures 3 and 4, and Table 1. Figure 3 shows the evolution and convergence of the FR metric as a function of the iteration number and Figure 4 shows the final ROS configuration. Table 1 contains the FR values of the original BPOF, the optimized TPAF, and for comparison purposes, ternary filters designed to maximize SNR[1], peak-to-correlation-energy PCE[11], and discrimination D defined as the ratio of the in-class peak intensity to the out-of-class peak intensity[12]. The latter three filters have analytic designs that maximize the metric of interest (SNR, PCE, or D) but of course are not optimized for FR. The TPAF designed with SA to maximize FR clearly provides the best FR performance, as expected. Transformation curves of FR to P_{ID} are given in Figure 5 for three different noise levels as a means of evaluating the designs in terms of P_{ID}.

3.2 Synthetic discriminant filters for distortion tolerance

As shown by the previous results, the SA design approach appears to work well for designing a ternary filter to maximize the probability of correct identification in a two image set. However, this is a very limited situation, and a much more useful filter would be one which can deal with larger numbers of in-class and out-of-class images. In particular, we are concerned with synthetic discriminant filters

(SDFs) which yield a degree of tolerance to distortions in the input image. For example, a commonly used type of SDF is that which produces equal correlation peak intensities for rotated versions of the basic in-class image. In this section, we address the design of this type of ternary filter that produces roughly equal peak intensities for the in-class image set, while maximizing the Fisher ratio defined by the in-class set and a similar set of rotated versions of the out-of-class image.

The synthetic discriminant function filter is designed to provide equal central correlation peak values for a given set of centered in-class training images. For a filter H(u,v), this is expressed as

$$\int S_j(u,v)\, H(u,v)\, dudv = k_j = k, \tag{15}$$

where $S_j(u,v)$ is the Fourier transform of the j-th training image $s_j(x,y)$, and k is a constant for all training images. A solution for the filter function H(u,v) involves creating a composite image $s_{comp}(x,y)$ as a linear combination of the training images as

$$s_{comp}(x,y) = \sum_{j=1}^{N_s} a_j s_j(x,y) \tag{16}$$

where N_s is the total number of in-class training images. The filter H(u,v) is then made from the composite image according to the modulation capabilities of the filter SLM as

$$H(u,v) = M\left\{S^*_{comp}(u,v)\right\} \tag{17}$$

where * represents the complex conjugate and $M\{X\}$ is the modulation operator. The objective of the SDF design algorithm is then to determine the coefficient vector $\mathbf{a} = [a_1\ a_2\ \ldots\ a_N]^T$ such that when the filter H(u) is made from $s_{comp}(x)$, the equal correlation peak condition expressed in Eq. 15 is satisfied.

For binary phase-only filters, the modulation operator is nonlinear and is represented by the binarization expressed in Eq. 7. In this case, there is no analytical solution for the proper coefficient vector **a** that yields equal correlation peaks for the training images. Instead, one must usually resort to iterative techniques for finding **a** such as the modified Newton-Raphson algorithm developed by Jared and Ennis[13] or the successive forcing algorithm of Bahri and Kumar[14]. Both of these algorithms essentially adjust the coefficient vector **a** to converge to a solution that satisfies Eq. 15. There is no theoretical guarantee that these algorithms will converge to a solution, but they have been used with reasonable success[13-15] and have converged in nearly all cases studied.

The general algorithm we employed to design ternary SDFs (TSDFs) that maximize a definition of the Fisher ratio was to first design a BSDF from the in-class training images using the technique of Jared and Ennis, and then use a simulated annealing procedure to design the ROS function M(u,v) that multiplies the BSDF. Thus the TSDF can be expressed as

$$TSDF(u,v) = M(u,v)\cdot BSDF(u,v) \tag{18}$$

and the design process is analogous to that used in the previous section with single in-class and out-of-class images. Now however, the definition of the Fisher ratio employed during the SA algorithm must be modified somewhat because the in-class and out-of-class sets both contain multiple images in general. If we make the reasonable assumption that each individual training image has an equal probability of being the input image correlated, then the logical modification of FR is to use the mean values of the in-class and out-of-class training image correlation peaks in the numerator. For the purpose of designing optimized TSDF filters, we then define the performance metric as

$$FR_{mean} = \frac{\left[\frac{1}{N_{ic}}\sum_{j=1}^{N_{ic}}|c_{ic,j}| - \frac{1}{N_{oc}}\sum_{j=1}^{N_{oc}}|c_{oc,j}|\right]^2}{\sigma_n^2 \sum_{u,v} M(u,v)} \qquad (19)$$

where N_{ic} and N_{oc} are the numbers of in-class and out-of-class training images, respectively.

Given the fact that the metric FR_{mean} as written in Eq. 19 is the measure of interest for a TSDF, it is obvious that this is the quantity we should maximize in the SA design procedure. However, as this requires evaluating the correlation peak values for N_{ic} in-class images and N_{oc} out-of-class images at every iteration in the algorithm, this design process will take much longer computationally in comparison to the design for single in-class and out-of-class images, especially for large N_{ic} and N_{oc}. Therefore we evaluate several different design procedures for the TSDF, and compare them in terms of their performance levels with respect to the metric FR_{mean} expressed in Eq. 19, and also in terms of the relative time required computationally to achieve the design. For the following four sets of results, we used 64x64 pixel rotated versions of the images shown in Figure 2. Six training images comprised both the in-class and out-of-class image sets, with the images rotated at every 1°, from 0° to 5°. This small range was chosen to decrease the design time, yet present a reasonably challenging design.

A. The first design process is simply the direct approach given above, where we maximize FR_{mean} during each iteration of the SA algorithm. As with the design in the previous section, we evaluate the in-class correlation peaks at the origin (since the in-class images are centered), and calculate the central 20x20 pixel region in the correlation plane for each out-of-class image, taking the maximum value in that region as the peak value. This process worked well, resulting in a ternary filter, designated $TSDF_a$, that has a metric value of $FR_{mean} = 5.2$. For comparison, the 5° binary SDF (BSDF) designed for these images has a value of $FR_{mean} = .83$.

B. Given the relatively long design time required by the direct method in A., we next tried to shorten the SA process by maximizing the Fisher ratio defined by the correlation peaks of the *composite* in-class and out-of-class images. That is, we defined $s_{ic,comp}(x,y) = 1/N_{ic}\Sigma s_{ic,j}(x,y)$, and similarly for $s_{oc,comp}(x,y)$, and then maximized the quantity

$$FR_{comp} = \frac{\left[|c_{ic,comp}| - |c_{oc,comp}|\right]^2}{\sigma_n^2 \sum_{u,v} M(u,v)} \qquad (20)$$

through design of the ROS function via simulated annealing. Note that $c_{ic,comp}$ and $c_{oc,comp}$ are the correlation peaks for the composite in-class image and composite out-of-class image, respectively. Thus this design is significantly faster than the direct approach because it involves evaluation of only a single in-class peak and a single out-of-class peak to compute FR_{comp} at each step. However, its disadvantage is that this algorithm does not directly maximize the quantity of interest. Accordingly, the result of this design after convergence was a filter $TSDF_b$ that has a metric value of $FR_{mean} = 4.0$, which is less than produced by $TSDF_a$, as we would expect.

C. Not surprisingly, the value of FR_{mean} is greater for $TSDF_a$ than for $TSDF_b$. However, the latter design also took far less time to compute, and it is in fact reasonably close in performance to the former. This suggests that perhaps a good design procedure might be to utilize both methods. Therefore in this approach, we effectively used a two-stage process by applying the direct but slow algorithm in A using the solution of the fast design in B as the starting point. The starting temperature of the second part (the direct optimization) was adjusted to be relatively low so that the starting point mask function was the foundation of the iterations. The result of this process was $TSDF_c$, which

produces a metric value of FR_{mean} = 5.1, which is nearly identical to the results found in A. However, the total design time in this case was significantly less than in A.

D. Finally, we can reduce the design time again by making an assumption regarding the in-class training image correlation peaks. Writing out the expression for the mean of the in-class peaks, we have

$$\frac{1}{N_{ic}}\sum_{j} c_{ic,j} = \frac{1}{N_{ic}}\sum_{j}\left|\sum_{u,v} S_j(u,v)\, H(u,v)\right| . \tag{21}$$

If we now assume that the in-class peak values are predominantly real and positive, then this can be re-written as

$$\begin{aligned}\frac{1}{N_{ic}}\sum_{j} c_{ic,j} &\cong \frac{1}{N_{ic}}\left|\sum_{j}\sum_{u,v} S_j(u,v)\, H(u,v)\right| \\ &= \left|\sum_{u,v}\left[\frac{1}{N_{ic}}\sum_{j} S_j(u,v)\right] H(u,v)\right|\end{aligned} \tag{22}$$

which upon inspection is seen to be the peak value of the composite in-class image. Thus if the assumption holds, then we may evaluate only the composite in-class image peak rather than all of the individual in-class image peaks. We can not make an analogous assumption for the out-of-class image peaks, however, because in this case the training image peak values may be anywhere in the correlation plane, and we must continue to evaluate each of the out-of-class correlation functions. Thus the method applied here is identical to C, except that in the second part of the optimization process, we continue to use the composite in-class image peak. This further decreases the total design time, while offering similar performance in terms of the mean Fisher ratio metric.

The results of the simulations performed for the SDF design algorithms A-D are summarized in Table 2. The results include not only FR_{mean}, but also a design time parameter t' that is basically the total number of iterations required multiplied by the time per iteration. This quantity is normalized to the direct approach outlined in A.

E. Since the design time using the simulated annealing algorithm can be quite long, a concern in the design of synthetic discriminant filters is how many training images, or equivalently, how large a distortion range, can be designed in this way. As we saw in section D, if we make the assumption that the in-class correlation peaks are predominantly real, then we can use the composite in-class image during the design which can significantly reduce the design length. Therefore, the real bottleneck is the number of out-of-class training images because we must generally evaluate each of their correlation functions at every iteration. Thus two filters that encompass distortion ranges that are different by a factor of K can be expected to require design times of roughly the same proportion. In this section we address the issue of designing larger range TSDFs without tremendously long design times. Towards that end, we performed design simulations for a 30° rotation range SDF, assuming that each of the rotated versions of the in-class and out-of-class images at every 1° were equally likely to appear as the input image. If we were to design this filter directly using every out-of-class training image, there would be a total of 31 such images, and the design process, even using the approach in section D, might be extremely long. Thus we reduced the number of out-of-class training images by using only those images spaced by a specific interval.

In the first approach, we used only those 4 out-of-class training images rotated at 0, 10, 20, and 30 degrees. The design was accomplished as in D, where we first made a preliminary design using the composite in-class and out-of-class images, and then refined the design by using the individual training images identified. A second filter was also designed using the 7 out-of-class training images rotated at 0, 5, 10, 15, 20, 25, and 30 degrees in an identical fashion, and finally a third filter was designed using the out-of-class training images spaced at every 3°. The results from the three designs are presented in Table 3. It is clear from those results that the 10° out-of-class image spacing is too large

to adequately represent the entire out-of-class set during the design process because the value of FR_{mean} obtained evaluating only the reduced set of 4 out-of-class training images does not correspond well to the value of FR_{mean} obtained when we evaluate all of the possible images in the range. However, as we decrease the spacing between the out-of-class images included in the design, the filter begins to be more representative of the entire out-of-class image set. The difference between the two values of FR_{mean} decreases progressively in the results for the 5° spacing design and the 3° spacing design. This approach can yield a good filter design in a small fraction of the time required if one uses all out-of-class training images. The necessary image spacing will depend of the particular image sets.

4. SUMMARY

We have studied the use of the simulated annealing optimization algorithm for the design of ternary phase and amplitude correlation filters for implementation on the magneto-optic spatial light modulator. Specifically, the SA technique was employed to maximize the Fisher ratio performance metric, which is directly related to the probability of correct identification between two classes in the presence of noise. We designed filters in this way for a single pair of images to be differentiated, and also for sets of in-class and out-of-class images corresponding to geometrically distorted versions of the base images. We obtained significant improvements in performance with filters designed with this algorithm in comparison to binary filters and other ternary filters optimized for criteria such as SNR and the discrimination ratio. The main drawback to this design approach is potentially long design times, so several methods to shorten the design were studied.

5. REFERENCES

1. B. V. K. Vijaya Kumar and Z. Bahri, "Efficient algorithm for designing a ternary valued filter yielding maximum signal to noise ratio," Appl. Opt. **28**, 1919-1925, (1989).
2. B. V. K. Vijaya Kumar, W. Shi, and C. Hendrix, "Phase-only filters with maximally sharp correlation peaks," Opt. Lett. **15**, 807-809, (1990).
3. J. D. Downie, "Design of optimal binary phase and amplitude filters for maximization of correlation peak sharpness," Opt. Lett. **16**, 508-510, (1991).
4. J. L. Horner, "Clarification of Horner efficiency," Appl. Opt. **31**, 4629-4629, (1992).
5. B. V. K. Vijaya Kumar and J. D. Brasher, "Relationship between maximizing the signal-to-noise ratio and minimizing the classification error probability for correlation filters," Opt. Lett. **17**, 940-942, (1992).
6. G. Zalman and J. Shamir, "Maximum discrimination filter," J. Opt. Soc. Am. A **8**, 814-821, (1991).
7. G. Zalman and J. Shamir, "Reducing error probability in optical pattern recognition," J. Opt. Soc. Am. A **8**, 1866-1873, (1991).
8. R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, Wiley, New York, (1973).
9. B. A. Kast, M. K. Giles, S. D. Lindell, and D. L. Flannery, "Implementation of tenary phase amplitude filters using a magnetooptic spatial light modulator," Appl. Opt. **28**, 1044-1046, (1989).
10. M. S. Kim, M. R. Feldman, and C. C. Guest, "Optimum encoding of binary phase-only filters with a simulated annealing algorithm," Opt. Lett. **14**, 545-547, (1989).
11. B. V. K. Vijaya Kumar, W. Shi, and C. Hendrix, "Partial information correlation filters with maximally sharp correlation peaks," Opt. Comp. and Proc. **1**, 29-46, (1991).
12. D. L. Flannery, J. S. Loomis, and M. E. Milkovich, "Transform-ratio ternary phase-amplitude filter formulation for improved correlation discrimination," Appl. Opt. **27**, 4079-4083, (1988).
13. D. A. Jared and D. J. Ennis, "Inclusion of filter modulation in synthetic-discriminant-function construction," Appl. Opt. **28**, 232-239, (1989).

14. Z. Bahri and B. V. K. Vijaya Kumar, "Binary phase-only synthetic discriminant functions (BPOSDFs) designed using the successive forcing algorithm," in *Hybrid Image and Signal Processing II*, Proc. SPIE **1297**, 188-193, (1990).
15. M. B. Reid, P.W. Ma, J. D. Downie, and E. Ochoa, "Experimental verification of modified synthetic discriminant function filters for rotation invariance," Appl. Opt. **29**, 1209-1214, (1990).

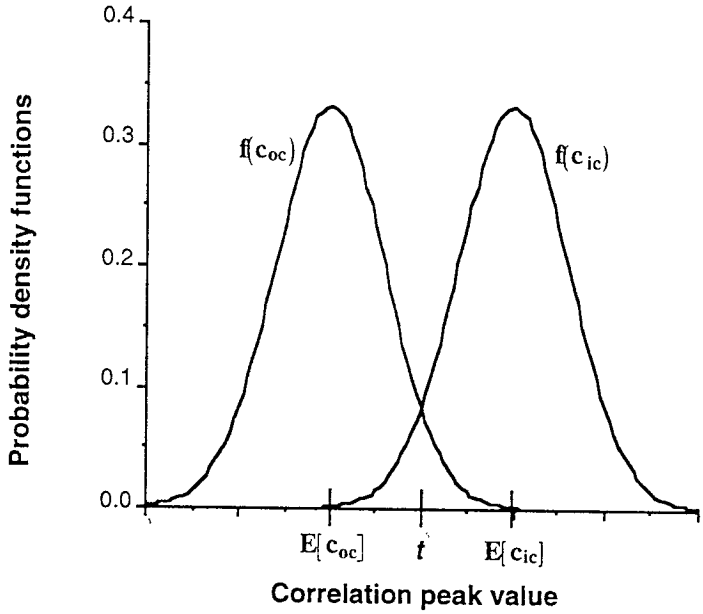

Figure 1: Gaussian distributions of in-class and out-of-class correlation peak values when input images suffer additive white noise. c_{ic} represents the in-class peak, c_{oc} represents the out-of-class peak, and t is the threshold value used to discriminate between the two classes. E[ ] represents the expected value.

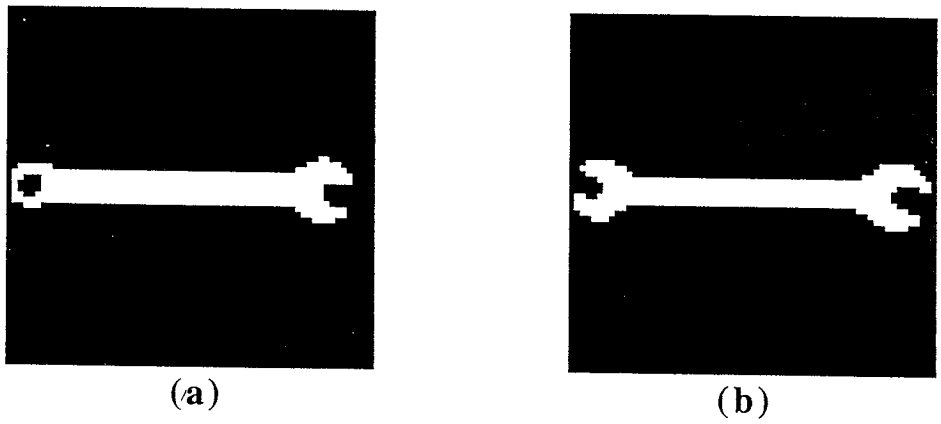

Figure 2: (a) 64x64 pixel in-class image. (b) 64x64 out-of-class image.

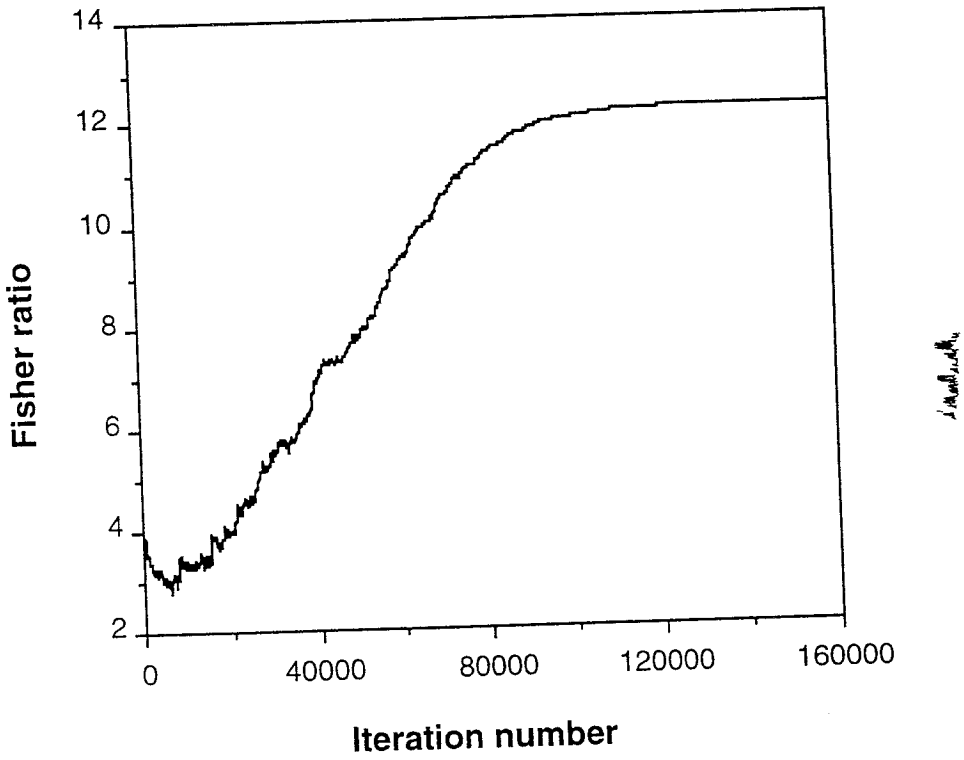

Figure 3: Evolution of FR as function of iteration number in simulated annealing algorithm for single in-class and out-of-class images.

Figure 4: Region of support (ROS) function for ternary filter that maximizes FR for single in-class and out-of-class images.

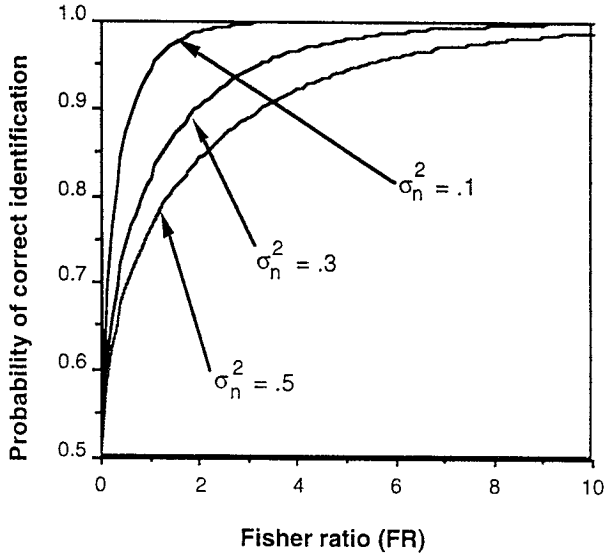

Figure 5: P_{ID} vs. Fisher ratio for three levels of image noise.

Filter	Fisher ratio FR
BPOF	3.9
$TPAF_{FR}$	12.2
$TPAF_{SNR}$	.3
$TPAF_{PCE}$	3.2
$TPAF_{D}$	2.8

Table 1: **$TSDF_{FR}$:** filter designed with simulated annealing to maximize FR, **$TSDF_{SNR}$:** filter designed to maximize SNR of in-class image, **$TSDF_{PCE}$:** filter designed to maximize PCE of in-class image, **$TSDF_{D}$:** filter designed to maximize discrimination ratio.

Filter	FR_{mean}	t'
BSDF	0.8	--
$TSDF_a$	5.2	1.00
$TSDF_b$	4.0	0.16
$TSDF_c$	5.1	0.45
$TSDF_d$	5.1	0.37

Table 2: Summary of results from TSDF design approaches A - D for a 5° range rotation-invariant filter. t' represents a relative measure of the total design time required.

Out-of-class image spacing	FR_{mean} (reduced set)	FR_{mean} (all training images)
10°	1.3	.6
5°	1.1	.8
3°	1.1	.9

Table 3: Three TSDF filters designed to span 30° rotation range.